

Policy Gradients for CVaR-Constrained MDPs

Prashanth L A*

INRIA Lille - Nord Europe, Team SequeL, FRANCE.

May 13, 2014

Abstract

We study a risk-constrained version of the stochastic shortest path (SSP) problem, where the risk measure considered is Conditional Value-at-Risk (CVaR). We propose two algorithms that obtain a locally risk-optimal policy by employing four tools: stochastic approximation, mini batches, policy gradients and importance sampling. Both the algorithms incorporate a CVaR estimation procedure, along the lines of Bardou et al. [2009], which in turn is based on Rockafellar-Uryasev's representation for CVaR and utilize the likelihood ratio principle for estimating the gradient of the sum of one cost function (objective of the SSP) and the gradient of the CVaR of the sum of another cost function (in the constraint of SSP). The algorithms differ in the manner in which they approximate the CVaR estimates/necessary gradients - the first algorithm uses stochastic approximation, while the second employ mini-batches in the spirit of Monte Carlo methods. We establish asymptotic convergence of both the algorithms. Further, since estimating CVaR is related to rare-event simulation, we incorporate an importance sampling based variance reduction scheme into our proposed algorithms.

1 Introduction

Risk-constrained Markov decision processes (MDPs) have attracted a lot of attention recently in the reinforcement learning (RL) community (cf. Borkar and Jain [2010], Tamar et al. [2012], Prashanth and Ghavamzadeh [2013], Tamar and Mannor [2013]). However, unlike previous works that focused mostly on variance of the return as a measure of risk, we consider Conditional Value-at-Risk (CVaR) as a risk measure. CVaR has the form of a conditional expectation, where the conditioning is based on a constraint on Value-at-Risk (VaR).

The aim in this paper is to find a *risk-optimal* policy in the context of a stochastic shortest path (SSP) problem. A risk-optimal policy is one that minimizes the sum of one cost function (see $G^\theta(s^0)$ in (1)), while ensuring that the conditional expectation of the sum of another cost function (see $C^\theta(s^0)$ in (1)) given some confidence level, stays bounded, i.e., the solution to the following risk-constrained problem: For a given $\alpha \in (0, 1)$ and $K_\alpha > 0$,

$$\min_{\theta \in \Theta} \mathbb{E} \left[\underbrace{\sum_{m=0}^{\tau-1} g(s_m, a_m) | s_0 = s^0}_{G^\theta(s^0)} \right] \text{ subject to } \text{CVaR}_\alpha \left[\underbrace{\sum_{m=0}^{\tau-1} c(s_m, a_m) | s_0 = s^0}_{C^\theta(s^0)} \right] \leq K_\alpha. \quad (1)$$

In the above, s^0 is the starting state and the actions $a_0, \dots, a_{\tau-1}$ are chosen according to a randomized policy π_θ governed by θ . Further, $g(s, a)$ and $c(s, a)$ are cost functions that take a state s and an action a as inputs and τ is the first passage time to the recurrent state of the underlying SSP (see Section 2 for a detailed formulation). In Borkar and Jain [2010], a similar problem is considered in a finite horizon MDP, though under a strong separability assumption for the cost function $c(s, a)$.

*prashanth.la@inria.fr

Solving the risk-constrained problem (1) is challenging due to two reasons:

- (i) Finding a globally risk-optimal policy is intractable even for a simpler case when the risk is defined as the variance of the return of an MDP (see Mannor and Tsitsiklis [2011]). The risk-constrained MDP that we consider is more complicated in comparison, since CVaR is a conditional expectation, with the conditioning governed by an event that bounds a probability.
- (ii) For the sake of optimization of the CVaR-constrained MDP that we consider in this paper, it is required to estimate both VaR/CVaR of the total cost ($C^\theta(s^0)$ in (1)) as well as its gradient. The problem is further complicated by the fact VaR/CVaR concerns the tail of the distribution of the total cost and hence, a variance reduction technique is required to speed up the estimation procedure.

We avoid the first problem by proposing a policy gradient scheme that is proven to converge to a locally optimal policy, while the second problem is alleviated using two principled approaches: stochastic approximation/mini-batch procedure for estimating VaR/CVaR derived out of a well-known convex optimization problem by Rockafellar and Uryasev [2000] and likelihood ratio estimates from the classic policy-gradient algorithm by Bartlett and Baxter [2011].

The contributions of this paper are summarized as follows:

(I) First, using the representation of CVaR (and also VaR) as the solution of a certain convex optimization problem by Rockafellar and Uryasev [2000], we develop a stochastic approximation procedure, along the lines of Bardou et al. [2009], for estimating the CVaR of a policy for an SSP. In addition, we also propose a scheme based on the mini-batch principle to estimate CVaR. Mini-batches are in the spirit of Monte Carlo methods and have been proposed by Atchade et al. [2014] under a different optimization context for stochastic proximal gradient algorithms.

(II) Second, we develop two novel policy gradient algorithms for finding a (locally) risk-optimal policy of the CVaR-constrained SSP. The first algorithm is a four time-scale stochastic approximation scheme while the second operates along two time-scales in conjunction with mini-batches. Both algorithms use a procedure to estimate CVaR and then use the policy-gradient principle with likelihood ratios to estimate the gradient of the total cost $G^\theta(s^0)$ as well as CVaR of another cost sum $C^\theta(s^0)$. Using the CVaR estimates as well as the necessary gradients (estimated along the fastest two time-scales), the first algorithm updates the policy parameter in the negative descent direction on the intermediate timescale and performs dual ascent for the Lagrange multiplier on the slowest timescale. On the other hand, the second algorithm operates on two timescales as it employs mini-batches to estimate the CVaR as well as the necessary gradients.

(III) Third, we adapt our proposed algorithms to incorporate importance sampling (IS) - a well-known variance-reduction scheme. This is motivated by the fact that when the confidence level α is close to 1, estimating VaR as well as CVaR takes longer. This is because the interesting samples used to estimate CVaR come from the tail of the distribution of the random variable concerned (in our case, the total cost $C^\theta(s^0)$) and thus, get rarer as α gets close to 1. Importance sampling (IS) is a standard tool to alleviate this problem and we employ the IS scheme proposed by Lemaire and Pages [2010]. However, applying the latter scheme in a SSP setting is non-trivial as it requires the knowledge of transition dynamics. We propose a heuristic IS variant where we use the randomized policies to derive sampling ratios for the IS procedure.

To sum up, the core contribution of this paper is twofold. First, using a careful synthesis of well-known techniques from stochastic approximation, likelihood ratios and importance sampling, we propose a policy gradient algorithm that is provably convergent to a locally risk-optimal policy. Second, we propose another algorithm based on the idea of mini-batches for estimating CVaR from the simulated samples. The latter approach is novel even for policy gradients in the context of risk-neutral MDPs.

The rest of the paper is organized as follows: In Section 2 we formalize the CVaR-constrained SSP and in Section 3 describe the structure of our proposed algorithms. In Section 4 we present the first algorithm based on stochastic approximation and in Section 5 we present the mini-batch variant. In Section 6, we sketch the convergence of our algorithms and later in Section 7 describe the importance sampling variants of our algorithms. In Section 8, we review relevant previous works. Finally, in Section 9 we provide the concluding remarks.

2 Problem formulation

In this section, we first introduce VaR/CVaR risk measures, then formalize the stochastic shortest path problem and subsequently define the CVaR-constrained SSP.

2.1 Background on VaR and CVaR

For any random variable X , we define the VaR at level $\alpha \in (0, 1)$ as

$$\text{VaR}_\alpha(X) := \inf \{ \xi \mid \mathbb{P}(X \leq \xi) \geq \alpha \}.$$

If the distribution of X is continuous, then VaR is the lowest solution to $\mathbb{P}(X \leq \xi) = \alpha$. VaR as a risk measure has several drawbacks, which precludes using standard stochastic optimization methods. This motivated the definition of coherent risk measures by Artzner et al. [1999]. A risk measure is coherent if it is convex, monotone, positive homogeneous and translation equi-variant. CVaR is one popular risk measure defined by

$$\text{CVaR}_\alpha(X) := \mathbb{E}[X \mid X \geq \text{VaR}_\alpha(X)].$$

Unlike VaR, the above is a coherent risk measure.

2.2 Stochastic Shortest Path (SSP)

We consider a SSP with a finite state space $\mathcal{S} = \{0, 1, \dots, r\}$, where 0 is a special cost-free terminal state. The set of feasible actions in state $s \in \mathcal{S}$ is denoted by $\mathcal{A}(s)$. A transition from state s to s' under action $a \in \mathcal{A}(s)$ occurs with probability $p_{ss'}(a)$ and incurs the following costs: $g(s, a)$ and $c(s, a)$, respectively. The terminal state 0 is cost-free and absorbing.

A policy specifies how actions are chosen in each state. A *stationary* randomized policy $\pi(\cdot|s)$ maps any state s to a probability vector on $\mathcal{A}(s)$. As is standard in policy gradient algorithms, we parameterize the policy and assume that the policy is continuously differentiable in the parameter θ . Since a policy π is identifiable by its parameter θ , we use them interchangeably in this paper.

As defined by Bertsekas [2007], a proper policy is one which ensures that there is a positive probability that the terminal state 0 will be reached, starting from any initial state, after utmost r transitions. This in turn implies the states $1, \dots, r$ are transient. We assume that class of parameterized policies considered, i.e., $\{\pi_\theta \mid \theta \in \Theta\}$, is proper.

2.3 CVaR-constrained SSP

As outlined earlier, the risk-constrained objective is:

$$\min_{\theta \in \Theta} \mathbb{E} \left[\underbrace{\sum_{m=0}^{\tau-1} g(s_m, a_m) \mid s_0 = s^0}_{G^\theta(s^0)} \right] \text{ subject to } \text{CVaR}_\alpha \left[\underbrace{\sum_{m=0}^{\tau-1} c(s_m, a_m) \mid s_0 = s^0}_{C^\theta(s^0)} \right] \leq K_\alpha,$$

where τ denotes the first visiting time to terminal state 0, i.e., $\tau = \min\{m \mid s_m = 0\}$. The actions $a_0, \dots, a_{\tau-1}$ are chosen according to the randomized policy π_θ . Further, α and K_α are constants that specify the confidence level and constraint bound for CVaR, respectively.

Using the standard trick of Lagrangian relaxation for constrained optimization problems, we convert (1) to the following unconstrained problem:

$$\max_{\lambda} \min_{\theta} [\mathcal{L}^{\theta, \lambda}(s^0) := G^\theta(s^0) + \lambda(\text{CVaR}_\alpha(C^\theta(s^0)) - K_\alpha)]. \quad (2)$$

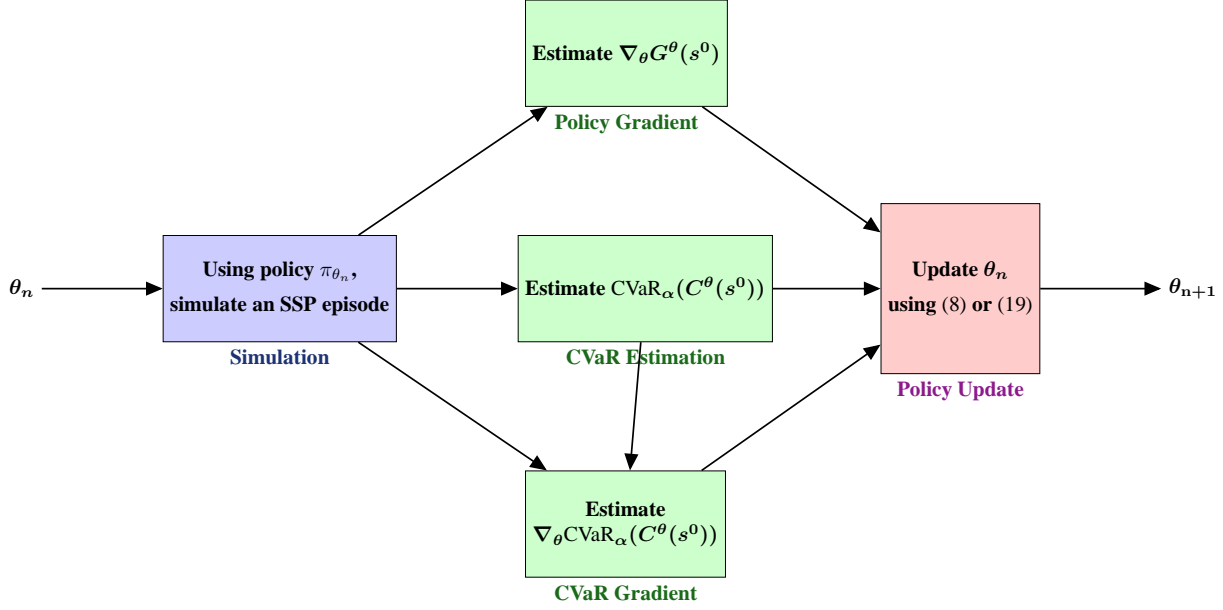


Figure 1: Overall flow of our algorithms.

3 Algorithm Structure

In order to solve (2), a standard constrained optimization procedure operates as follows:

Simulation. This is the inner-most loop where the SSP is simulated for several episodes and the resulting costs are aggregated.

Policy Update. This is the intermediate loop where the gradient of the Lagrangian along θ is estimated using simulated values above. The gradient estimates are then used to update policy parameter θ along a descent direction. Note that this loop is for a given value of λ ; and

Lagrange Multiplier Update. This is the outer-most loop where the Lagrange multiplier λ is updated along an ascent direction, using the converged values of the inner two loops.

Using two-timescale stochastic approximation (see Chapter 6 of Borkar [2008]), the policy and Lagrange multiplier update can run in parallel as follows:

$$\theta_{n+1} = \theta_n - \gamma_n \nabla_{\theta} \mathcal{L}^{\theta, \lambda}(s^0) \quad \text{and} \quad \lambda_{n+1} = \Gamma_{\lambda} [\lambda_n + \beta_n \nabla_{\lambda} \mathcal{L}^{\theta, \lambda}(s^0)], \quad (3)$$

where Γ_{λ} is a projection operator that keep the iterate λ_n bounded, while $\gamma_n, \beta_n, n \geq 0$ are step-sizes that satisfy the following assumption:

$$\sum_{n=1}^{\infty} \beta_n = \infty, \quad \sum_{n=1}^{\infty} \gamma_n = \infty, \quad \sum_{n=1}^{\infty} (\beta_n^2 + \gamma_n^2) < \infty \quad \text{and} \quad \frac{\beta_n}{\gamma_n} \rightarrow 0.$$

The last condition above ensures that θ -recursion proceeds on a faster timescale in comparison to λ -recursion.

Simulation optimization. No closed form expression for the gradient of the Lagrangian $L^{\theta, \lambda}(s^0)$ is available and moreover, $G^{\theta}(s^0)$ and $C^{\theta}(s^0)$ are observable only via simulation. Observe that $\nabla_{\theta} \mathcal{L}^{\theta, \lambda}(s^0) = \nabla_{\theta} G^{\theta}(s^0) + \lambda \nabla_{\theta} \text{CVaR}_{\alpha}(C^{\theta}(s^0))$ and $\nabla_{\lambda} \mathcal{L}^{\theta, \lambda}(s^0) = \text{CVaR}_{\alpha}(C^{\theta}(s^0)) - K_{\alpha}$. Hence, in order to update according to (3), we need to estimate, for any policy parameter θ , the following quantities via simulation:

(i) $\text{CVaR}_{\alpha}(C^{\theta}(s^0))$; (ii) $\nabla_{\theta} G^{\theta}(s^0)$; and (iii) $\nabla_{\theta} \text{CVaR}_{\alpha}(C^{\theta}(s^0))$.

In the following sections, we describe two algorithms that differ in the way they estimate each of the above quantities and subsequently establish that the estimates (and hence the overall algorithms) converge.

4 Algorithm 1: PG-CVaR-SA

Algorithm 1 describes the complete algorithm along with the update rules for the various parameters. The algorithm involves the following crucial components - simulation of the SSP, VaR/CVaR estimation and policy gradients for the objective as well as the CVaR constraint. Each of these components is described in detail in the following.

Algorithm 1 PG-CVaR-SA

Input: parameterized policy $\pi_\theta(\cdot|\cdot)$, step-sizes $\{\zeta_{n,1}, \zeta_{n,2}, \gamma_n, \beta_n\}_{n \geq 1}$
Initialization: Starting state s^0 , initial policy θ_0 , , number of iterations $M \gg 1$.
for $n = 1, 2, \dots, M$ **do**
 Simulation

 Simulate the SSP for an episode using actions $a_{n,0}, \dots, a_{n,\tau_n-1}$ generated using $\pi_{\theta_{n-1}}$

 Obtain cost estimates: $G_n := \sum_{j=0}^{\tau_n-1} g(s_{n,j}, a_{n,j})$ and $C_n := \sum_{j=0}^{\tau_n-1} c(s_{n,j}, a_{n,j})$

 Obtain likelihood derivative: $z_n := \sum_{j=0}^{\tau_n-1} \nabla \log \pi_\theta(s_{n,j}, a_{n,j})$

VaR/CVaR estimation:

$$\text{VaR: } \xi_n = \xi_{n-1} - \zeta_{n,1} \left(1 - \frac{1}{1-\alpha} \mathbf{1}_{\{C_n \geq \xi_{n-1}\}} \right), \quad (4)$$

$$\text{CVaR: } \psi_n = \psi_{n-1} - \zeta_{n,2} (\psi_{n-1} - v(\xi_{n-1}, C_n)). \quad (5)$$

Policy Gradient:

$$\text{Total Cost: } \bar{G}_n = \bar{G}_{n-1} - \gamma_n (G_n - \bar{G}_n), \quad \text{Gradient: } \partial G_n = \bar{G}_n z_n. \quad (6)$$

CVaR Gradient:

$$\text{Total Cost: } \tilde{C}_n = \tilde{C}_{n-1} - \gamma_n (C_n - \tilde{C}_n), \quad \text{Gradient: } \partial C_n = (\tilde{C}_n - \xi_n) z_n \mathbf{1}_{\{C_n \geq \xi_n\}}. \quad (7)$$

Policy and Lagrange Multiplier Update:

$$\theta_n = \theta_{n-1} - \gamma_n (\partial G_n + \lambda_{n-1} (\partial C_n)), \quad \lambda_n = \Gamma_\lambda (\lambda_{n-1} + \beta_n (\psi_n - K_\alpha)). \quad (8)$$

end for

Output (θ_M, λ_M) .

4.1 SSP Simulation

In each iteration of PG-CVaR-SA, an episode of the underlying SSP is simulated. Each episode ends with a visit to the recurrent state 0 of the SSP. Let τ_n denote the time of this visit in episode n . The actions $a_{n,j}, j = 0, \dots, \tau_n - 1$ in episode n are chosen according to the policy $\pi_{\theta_{n-1}}$. Let $G_n := \sum_{j=0}^{\tau_n-1} g(s_{n,j}, a_{n,j})$ and $C_n := \sum_{j=0}^{\tau_n-1} c(s_{n,j}, a_{n,j})$ denote the accumulated cost values. Further, let $z_n := \sum_{j=0}^{\tau_n-1} \nabla \log \pi_\theta(s_{n,j}, a_{n,j})$ denote the likelihood derivative

(see Section 4.3 below). The tuple (G_n, C_n, z_n) obtained at the end of the n th episode is used to estimate CVaR as well as policy gradients.

4.2 Estimating VaR and CVaR

A well-known result by Rockafellar and Uryasev [2000] is that both VaR and CVaR can be obtained from the solution of a certain convex optimization problem and we recall this result next.

Theorem 1 *For any random variable X , let*

$$v(\xi, X) := \xi + \frac{1}{1-\alpha}(X - \xi)_+ \text{ and } V(\xi) = \mathbb{E}[v(\xi, X)] \quad (9)$$

Then, $\text{VaR}_\alpha(X)$ is any point of the set $\arg \min V = \{\xi \in \mathbb{R} \mid V'(\xi) = 0\}$, where V' is the derivative of V w.r.t. ξ . Further, $\text{CVaR}_\alpha(X) = V(\xi_\alpha^)$, where ξ_α^* is a $\text{VaR}_\alpha(X)$.*

From the above, it is clear that in order to estimate VaR/CVaR, one needs to find a ξ that satisfies $V'(\xi) = 0$. Stochastic approximation (SA) is a natural tool to use in this situation. We briefly introduce SA next and later develop a scheme for estimating CVaR along the lines of Bardou et al. [2009] on the faster timescale of PG-CVaR-SA.

Stochastic approximation. The aim is to solve the equation $F(\theta) = 0$ when analytical form of F is not known. However, noisy measurements $F(\theta_n) + \xi_n$ can be obtained, where $\theta_n, n \geq 0$ are the input parameters and $\xi_n, n \geq 0$ are zero-mean random variables, that are not necessarily i.i.d.

The seminal Robbins Monro algorithm solved this problem by employing the following update rule:

$$\theta_{n+1} = \theta_n + \gamma_n(F(\theta_n) + \xi_n). \quad (10)$$

In the above, γ_n are step-sizes that satisfy $\sum_{n=1}^{\infty} \gamma_n = \infty$ and $\sum_{n=1}^{\infty} \gamma_n^2 < \infty$. Under a stability assumption for the iterates and bounded noise, it can be shown that θ_n governed by (10) converges to the solution of $F(\theta) = 0$ (cf. Proposition 1 in Section 6).

4.2.1 CVaR estimation using SA.

Using the stochastic approximation principle and the result in Theorem 1, we have the following scheme to estimate the VaR/CVaR simultaneously from the simulated samples C_n :

$$\text{VaR: } \xi_n = \xi_{n-1} - \underbrace{\zeta_{n,1} \left(1 - \frac{1}{1-\alpha} \mathbf{1}_{\{C_n \geq \xi\}}\right)}_{\frac{\partial v}{\partial \xi}(\xi, C_n)}, \quad (11)$$

$$\text{CVaR: } \psi_n = \psi_{n-1} - \zeta_{n,2} (\psi_{n-1} - v(\xi_{n-1}, C_n)). \quad (12)$$

In the above, (11) can be seen as a gradient descent rule, while (12) can be seen as a plain averaging update. The scheme above is similar to the one proposed by Bardou et al. [2009], except that the random variable $C^\theta(s^0)$ (whose CVaR we try to estimate) is a sum of costs obtained at the end of each episode, unlike the single-shot r.v. considered by Bardou et al. [2009]. The step-sizes $\zeta_{n,1}, \zeta_{n,2}$ satisfy

$$\sum_{n=1}^{\infty} \zeta_{n,1} = \infty, \sum_{n=1}^{\infty} \zeta_{n,2} = \infty, \sum_{n=1}^{\infty} (\zeta_{n,1}^2 + \zeta_{n,2}^2) < \infty, \frac{\zeta_{n,2}}{\zeta_{n,1}} \rightarrow 0 \text{ and } \frac{\gamma_n}{\zeta_{n,2}} \rightarrow 0.$$

The last two conditions above ensure that VaR estimation recursion (11) proceeds on a faster timescale in comparison to CVaR estimation recursion (12) and further, the CVaR recursion itself proceeds on a faster timescale as compared to the policy parameter θ -recursion.

Using the ordinary differential equation (ODE) approach, we establish later that the tuple (ξ_n, ψ_n) converges to $\text{VaR}_\alpha(C^\theta(s^0)), \text{CVaR}_\alpha(C^\theta(s^0))$, for any fixed policy parameter θ (see Section 6).

4.3 Policy Gradient

We briefly introduce the technique of likelihood ratios for gradient estimation Glynn [1987] and later provide the necessary estimate for the gradient of total cost $G^\theta(s^0)$.

4.3.1 Gradient estimation using likelihood ratios.

Consider a Markov chain $\{X_n\}$ with a single recurrent state 0 and transient states $1, \dots, r$. Let $P(\theta) := [[p_{X_i X_j}(\theta)]]_{i,j=0}^r$ denote the transition probability matrix of this chain. Here $p_{X_i X_j}(\theta)$ denotes the probability of going from state X_i to X_j and is parameterized by θ . Let τ denote the first passage time to the recurrent state 0.

Let $X := (X_0, \dots, X_{\tau-1})^T$ denote the sequence of states encountered between visits to the recurrent state 0. The aim is to optimize a performance measure $F(\theta) = \mathbb{E}[f(\theta, X)]$ for this chain using simulated values of X . The likelihood estimate is obtained by first simulating the Markov chain according to $P(\theta)$ to obtain the samples $X_0, \dots, X_{\tau-1}$ and then estimate the gradient as follows:

$$\nabla_\theta F(\theta) = \mathbb{E} \left[f(X) \sum_{m=0}^{\tau-1} \frac{\nabla_\theta p_{X_m X_{m+1}}(\theta)}{p_{X_m X_{m+1}}(\theta)} \right].$$

4.3.2 Policy Gradient for the objective.

For estimating the gradient of the objective $G^\theta(s^0)$, we employ the following well-known estimate (cf. Bartlett and Baxter [2011]):

$$\nabla_\theta G^\theta(s^0) = \mathbb{E} \left[\left(\sum_{n=0}^{\tau-1} g(s_n, a_n) \right) \nabla \log P(s_0, \dots, s_{\tau-1}) \mid s_0 = s^0 \right], \quad (13)$$

where $\nabla \log P(s_0, \dots, s_\tau)$ is the likelihood derivative for a policy parameterized by θ , defined as

$$\nabla \log P(s_0, \dots, s_{\tau-1}) = \sum_{m=0}^{\tau-1} \nabla \log \pi_\theta(a_m \mid s_m). \quad (14)$$

The above relation holds owing to the fact that we parameterize the policies and hence, the gradient of the transition probabilities can be estimated from the policy alone. This is the well-known policy gradient technique Bartlett and Baxter [2011] that makes it amenable for estimating gradient of a performance measure in MDPs, since the transition probabilities are not required and one can work with policies and simulated transitions from the MDP.

4.4 Policy Gradient for the CVaR constraint.

For estimating the gradient of the CVaR of $C^\theta(s^0)$ for a given policy parameter θ , we employ the following likelihood estimate proposed by Tamar et al. [2014]:

$$\begin{aligned} \nabla_\theta \text{CVaR}_\alpha(C^\theta(s^0)) \\ = \mathbb{E} \left[(C^\theta(s^0) - \text{VaR}_\alpha(C^\theta(s^0))) \nabla \log P(s_0, \dots, s_{\tau-1}) \mid C^\theta(s^0) \geq \text{VaR}_\alpha(C^\theta(s^0)) \right], \end{aligned} \quad (15)$$

where $\nabla \log P(s_0, \dots, s_\tau)$ is as defined before in (14).

Since we do not know $\text{VaR}_\alpha(C^\theta(s^0))$, in Algorithm 1 we have an online scheme that uses ξ_n (see (11)) to approximate $\text{VaR}_\alpha(C^\theta(s^0))$, which is then used to derive an approximation to the gradient $\nabla_\theta \text{CVaR}_\alpha(C^\theta(s^0))$ (see (7)).




Figure 2: Illustration of mini-batch principle in PG-CVaR-mB algorithm.

5 Algorithm 2: PG-CVaR-mB

mini-Batches. As illustrated in Figure 2, we simulate the SSP for several episodes in each iteration of PG-CVaR-mB. At the end of the simulation, we obtain the total costs and likelihood derivative estimates $\{G_{n,j}, C_{n,j}, z_{n,j}\}_{j=1}^{m_n}$ and using these, the following quantities are approximated: $\text{CVaR}_\alpha(C^\theta(s^0))$, $\nabla_\theta \text{CVaR}_\alpha(C^\theta(s^0))$ and $\nabla_\theta G^\theta(s^0)$. The latter approximations can be seen as empirical means of functions of $G_{n,j}, C_{n,j}, z_{n,j}$. The complete algorithm along with the update rules for various parameters is presented in Algorithm 2.

mini-Batch size. A simple setting for the batch-size m_n is Cn^δ for some $\delta > 0$, i.e., m_n increases as a function of n . We cannot have constant batches, i.e., $\delta = 0$, since the bias of the CVaR estimates and the gradient approximations has to vanish asymptotically.

6 Outline of Convergence

We analyze our algorithms using the theory of multiple time-scale stochastic approximation [Borkar, 2008, Chapter 6]. Both the algorithms comprise of updates to the policy parameter θ on the faster time-scale and to the Lagrange multiplier λ on the slower time-scale. We first provide the analysis for PG-CVaR-SA algorithm and later describe the necessary modification for the mini-batch variant.

Before the main proof, we recall the following well-known result (cf. Chapter 2 of Borkar [2008]) related to convergence of stochastic approximation schemes under the existence of a so-called *Lyapunov function*:

Proposition 1 *Consider the following recursive scheme:*

$$\theta_{n+1} = \theta_n + \gamma_n(F(\theta_n) + \xi_{n+1}), \quad (20)$$

where $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a L -Lipschitz map and ξ_n a square-integrable martingale difference sequence with respect to the filtration $\mathcal{F}_n := \sigma(\theta_m, \xi_m, m \leq n)$. Moreover, $\mathbb{E}[\|\xi_{n+1}\|_2^2 | \mathcal{F}_n] \leq K(1 + \|\theta_n\|_2^2)$ for some $K > 0$. The step-sizes γ_n satisfy $\sum_{n=1}^{\infty} \gamma_n = \infty$ and $\sum_{n=1}^{\infty} \gamma_n^2 < \infty$.

Lyapunov function. Suppose there exists a continuously differentiable $V : \mathbb{R}^d \rightarrow [0, \infty)$ such that $\lim_{\|\theta\|_2 \rightarrow \infty} V(\theta) = \infty$. Writing $\mathbb{Z} := \{\theta \in \mathbb{R}^d \mid V(\theta) = 0\} \neq \emptyset$, V satisfies $\langle F(\theta), \nabla V(\theta) \rangle \leq 0$ with equality if and only if $\theta \in \mathbb{Z}$.

Then, θ_n governed by (20) converges a.s. to an internally chain transitive set contained in \mathbb{Z} .

The steps involved in proving the convergence of PG-CVaR-SA are as follows:

Step 1: CVaR estimation on fastest time-scale

Owing to the time-scale separation, θ and λ can be assumed to be constant (quasi-static) while analyzing the VaR/CVaR estimation procedure. We first show that the VaR estimate ξ_n converges to the corresponding true value $\text{VaR}_\alpha(C^\theta(s^0))$. This can be inferred by observing that V (see (9)) itself serves as the Lyapunov function and the fact that the step-sizes satisfy (A3) implies the iterates remain bounded. Thus, by an application of Proposition 1, it is evident that the recursion (11) converges to a point in the set $\{\xi \mid V(\xi) = 0\}$. Since every local minimum is a global minimum for V , the iterates ξ_n will converge to $\text{VaR}_\alpha(C^\theta(s^0))$. Establishing the convergence of the companion recursion in (12) to estimate $\text{CVaR}_\alpha(C^\theta(s^0))$ is easier. This is because (12) is a plain averaging update that uses the VaR estimate ξ_n from (11).

Algorithm 2 PG-CVaR-mB

Input: parameterized policy $\pi_\theta(\cdot|\cdot)$, step-sizes $\{\gamma_n, \beta_n\}$, non-negative weights $\{a_n\}$, mini-batch sizes $\{m_n\}$.

Initialization: Starting state s^0 , initial policy θ_0 , number of iterations $M \gg 1$.

for $n = 1, 2, \dots, M$ **do**

Simulation

 Simulate the SSP for m_n episodes using actions generated from the policy $\pi_{\theta_{n-1}}$

 Obtain cost estimates: $\{G_{n,j}, C_{n,j}\}_{j=1}^{m_n}$

 Obtain likelihood derivatives: $\{z_{n,j}\}_{j=1}^{m_n}$

VaR/CVaR estimation:

$$\text{VaR: } \xi_n = \frac{1}{m_n} \sum_{j=1}^{m_n} \left(1 - \frac{\mathbf{1}_{\{C_{n,j} \geq \xi_{n-1}\}}}{1 - \alpha} \right), \quad \text{CVaR: } \psi_n = \frac{1}{m_n} \sum_{j=1}^{m_n} v(\xi_{n-1}, C_{n,j}) \quad (16)$$

Policy Gradient:

$$\text{Total Cost: } \bar{G}_n = \frac{1}{m_n} \sum_{j=1}^{m_n} G_{n,j}, \quad \text{Gradient: } \partial G_n = \bar{G}_n z_n. \quad (17)$$

CVaR Gradient:

$$\text{Total Cost: } \bar{C}_n = \frac{1}{m_n} \sum_{j=1}^{m_n} C_{n,j}, \quad \text{Gradient: } \partial C_n = (\tilde{C}_n - \xi_n) z_n \mathbf{1}_{\{\bar{C}_n \geq \xi_n\}}. \quad (18)$$

Policy and Lagrange Multiplier Update:

$$\theta_n = \theta_{n-1} - \gamma_n (\partial G_n + \lambda_{n-1} (\partial C_n)), \quad \lambda_n = \Gamma_\lambda (\lambda_{n-1} + \beta_n (\psi_n - K_\alpha)) \quad (19)$$

end for

$$\text{Output } \left(\bar{\theta}_M := \frac{\sum_{k=1}^M a_k \theta_k}{\sum_{k=1}^M a_k}, \lambda_M \right).$$

Step 2: Policy update on intermediate time-scale

We provide the main arguments to show that θ_t governed by (8) converges to asymptotically stable equilibrium points of the following ODE:

$$\dot{\theta}_t = \nabla_{\theta} \mathcal{L}^{\theta_t, \lambda}(s^0) = \nabla_{\theta} G^{\theta_t}(s^0) + \lambda \nabla_{\theta} \text{CVaR}_{\alpha}(C^{\theta_t}(s^0)) \quad (21)$$

Since λ is on the slowest timescale, its effect is 'quasi-static' on the θ -recursion. Further, since the CVaR estimation and necessary gradient estimates using likelihood ratios are on the faster timescale, the θ -update in (8) views these quantities as almost equilibrated. Thus, the θ -update in (8) can be seen to be asymptotically equivalent to the following in the sense that the difference between the two updates is $o(1)$:

$$\theta_{t+1} = \theta_t - \gamma_t (\nabla_{\theta} G^{\theta_t}(s^0) + \lambda \nabla_{\theta} \text{CVaR}_{\alpha}(C^{\theta_t}(s^0))) ,$$

Thus, (8) can be seen to be a discretization of the ODE (21). Moreover, $\mathcal{L}^{\theta, \lambda}(s^0)$ serves as the Lyapunov function for the above recursion, since $\frac{d\mathcal{L}^{\theta, \lambda}(s^0)}{dt} = \nabla_{\theta} \mathcal{L}^{\theta, \lambda}(s^0) \dot{\theta} = \nabla_{\theta} \mathcal{L}^{\theta, \lambda}(s^0) (-\nabla_{\theta} \mathcal{L}^{\theta, \lambda}(s^0)) < 0$. Thus, by an application of Kushner-Clark lemma Kushner and Clark [1978], θ -recursion in (8) can be seen to converge to the asymptotically stable attractor for the ODE (21).

Step 3: Lagrange multiplier update on slowest time-scale

This is easier in comparison to the other steps and follows using arguments similar to that used for constrained MDPs in general by Borkar [2005]. The λ recursion views θ as almost equilibrated owing to time-scale separation and converges to the set of asymptotically stable equilibria of the following system of ODEs:

$$\dot{\lambda}_t = \check{\Gamma}_{\lambda}(\nabla_{\lambda} \mathcal{L}^{\theta^{\lambda_t}, \lambda_t}(s^0)) = \check{\Gamma}_{\lambda}(\text{CVaR}_{\alpha}(C^{\theta^{\lambda_t}}(s^0)) - K_{\alpha}) \quad (22)$$

where θ^{λ} is the value of the converged policy parameter θ when multiplier λ is used. $\check{\Gamma}_{\lambda}$ is a suitably defined projection operator that keeps λ_t evolving according to (22) bounded. Next, the PG-CVaR-SA algorithm converges to the a (local) saddle point of $\mathcal{L}^{\theta, \lambda}(s^0)$, i.e., to a tuple (θ^*, λ^*) that are a local minimum w.r.t. θ and a local maximum w.r.t. λ of $\mathcal{L}^{\theta, \lambda}(s^0)$.

The two claims above related to the convergence of λ -recursion and overall convergence follow using arguments similar to that by Borkar [2005], Prashanth and Ghavamzadeh [2013]. In particular, the former claim follows using standard stochastic approximation by viewing λ -recursion as performing gradient ascent, whereas the latter claim requires invocation of the envelope theorem of mathematical economics Mas-Colell et al. [1995].

PG-CVaR-mB. The proof for mini-batch variant differs only in the first step, i.e., estimation of VaR/CVaR and necessary gradients. Assume that the number m_n of mini-batch samples used for averaging in (16)–(18), increases with n . Thus, a straightforward application of law of large numbers establishes that the empirical mean estimates in (16)–(18) converge to their corresponding true values. The rest of the proof follows in a similar manner as PG-CVaR-SA.

7 Extension to incorporate Importance Sampling

In this section, we incorporate an importance sampling procedure in the spirit of Lemaire and Pages [2010], Bardou et al. [2009] to speed up the estimation procedure for VaR/CVaR in our algorithms.

7.1 Importance sampling.

Given a random variable X with density $p(\cdot)$ and a function $H(\cdot)$, the aim of an IS based scheme is to estimate the expectation $\mathbb{E}(H(X))$ in a manner that reduces the variance of the estimates. Suppose X is sampled using

another distribution with density $\tilde{p}(X, \eta)$ (parameterized by η), such that $\tilde{p}(X, \eta) = 0 \Rightarrow p(X) = 0$ (an absolute continuity condition). Then,

$$\mathbb{E}(H(X)) = \mathbb{E} \left[H(X) \frac{p(X)}{\tilde{p}(X, \eta)} \right]. \quad (23)$$

The problem is to choose the parameter η of the sampling distribution so as to minimize the variance of the above estimate.

A slightly different approach based on mean-translation is taken in a recent method proposed by Lemaire and Pages [2010]. By translation invariance, we have

$$\mathbb{E}[H(X)] = \mathbb{E} \left[H(X + \eta) \frac{p(X + \eta)}{p(X)} \right], \quad (24)$$

and the objective is to find a η that minimizes the following variance term:

$$Q(\eta) := \mathbb{E} \left[H^2(X + \eta) \frac{p^2(X + \eta)}{p^2(X)} \right]. \quad (25)$$

If ∇Q can be written as an expectation, i.e., $\nabla Q(\eta) = \mathbb{E}[q(\eta, X)]$, then one can hope to estimate this expectation (and hence minimize Q) using a stochastic approximation recursion. However, this is not straightforward since $\|q(\eta, x)\|_2$ is required to be sub-linear to ensure convergence of the resulting scheme¹.

One can get around this problem by double translation of η as suggested first by Lemaire and Pages [2010] and later used by Bardou et al. [2009] for VaR/CVaR estimation. Formally, under classic log-concavity assumptions on $p(X)$, it can be shown that Q is finite, convex and differentiable, so that

$$\nabla Q(\eta) := \mathbb{E} \left[H(X - \eta)^2 \frac{p^2(X - \eta)}{p(X)p(X - 2\eta)} \frac{\nabla p(X - 2\eta)}{p(X - 2\eta)} \right]. \quad (26)$$

Writing $K(\eta, X) := \frac{p^2(X - \eta)}{p(X)p(X - 2\eta)} \frac{\nabla p(X - 2\eta)}{p(X - 2\eta)}$, one can bound $K(\eta, X)$ by a deterministic function of η as follows: $|K(\eta, X)| \leq e^{2\rho|\eta|^b} (A|x|^{b-1} + A|\eta|^{b-1} + B)$, for some constants ρ, A and B . The last piece before present an IS scheme is related to controlling the growth of $H(X)$. We assume that $H(X)$ is controlled by another function $W(X)$ that satisfies $\forall x, |H(x)| \leq W(x)$, $W(x + y) \leq C(1 + W(x))^c(1 + W(y))^c$ and $\mathbb{E}[|X|^{2(b-1)}W(X)^{4c}] < \infty$.

An IS scheme based on stochastic approximation updates as follows:

$$\eta_n = \eta_{n-1} - \gamma_n \tilde{q}(\eta_{n-1}, X_n), \quad (27)$$

where $\tilde{q}(\eta, X) := H(X - \eta)^2 e^{-2\rho|\eta|^b} K(\eta, X)$. In lieu of the above discussion, $\|\tilde{q}(\eta, X)\|_2$ can be bounded by a linear function of $\|\eta\|_2$ and hence, the recursion (27) converges to the set $\{\eta \mid \nabla Q(\eta) = 0\}$ (See Section 2.3 by Bardou et al. [2009] for more details).

7.2 IS for VaR/CVaR estimation.

Let $D := (s_0, a_0, \dots, s_{T-1}, a_{T-1})$ be the random variable corresponding to an SSP episode and let $D_n := (s_{n,0}, a_{n,0}, \dots, s_{n,T-1}, a_{n,T-1})$ be the n th sample simulated using the distribution of D . Recall that the objective is to estimate the VaR/CVaR of the total cost $C^\theta(s^0)$, for a given policy parameter θ using samples from D .

Applying the IS procedure described above to our setting is not straightforward, as one requires the knowledge of the density, say $p(\cdot)$, of the random variable D . Notice that the density $p(D)$ can be written as $p(D) =$

¹As illustrated by [Bardou et al., 2009, Section 2.3], even for a standard Gaussian distributed X , i.e., $X \sim \mathcal{N}(0, 1)$, the function $q(\eta, x) = \exp(|\eta|^2/2 - \eta x) H^2(x)(\eta - x)$ and hence not sub-linear.

$\prod_{m=0}^{\tau-1} \pi_\theta(a_m | s_m) P(s_{m+1} | s_m, a_m)$. As pointed out in earlier works (cf. Sutton and Barto [1998]), the ratio $\frac{p(d)}{p(d')}$ can be computed for two (independent) episodes d and d' without requiring knowledge of the transition dynamics.

In the following, we use $\tilde{p}(D_n) := \prod_{m=0}^{\tau-1} \pi_\theta(a_{n,m} | s_{n,m})$ as a proxy for the density $p(D_n)$ and apply the IS scheme described above to reduce the variance of the VaR/CVaR estimation scheme (11)–(12). The update rule of the resulting recursion is as follows:

$$\xi_n = \xi_{n-1} - \zeta_{n,1} e^{-\rho|\eta|^\beta} \left(1 - \frac{1}{1-\alpha} \mathbf{1}_{\{C_n + \eta_{n-1} \geq \xi_{n-1}\}} \frac{\tilde{p}(D_n + \eta_{n-1})}{\tilde{p}(D_n)} \right), \quad (28)$$

$$\eta_n = \eta_{n-1} - \zeta_{n,1} e^{-2\rho|\eta_{n-1}|^\beta} \mathbf{1}_{\{C_n - \eta_{n-1} \geq \xi_{n-1}\}} \frac{\tilde{p}^2(D_n - \eta_{n-1})}{\tilde{p}(D_n)\tilde{p}(D_n - 2\eta)} \frac{\nabla \tilde{p}(D_n - 2\eta_{n-1})}{\tilde{p}(D_n - 2\eta_{n-1})}. \quad (29)$$

$$\psi_n = \psi_{n-1} - \zeta_{n,2} \left(\psi_{n-1} - \xi_{n-1} - \frac{1}{1-\alpha} (C_n + \mu_{n-1} - \xi_{n-1}) \mathbf{1}_{\{C_n + \mu_{n-1} \geq \xi_{n-1}\}} \frac{\tilde{p}(D_n + \mu_{n-1})}{\tilde{p}(D_n)} \right), \quad (30)$$

$$\begin{aligned} \mu_n = \mu_{n-1} - \zeta_{n,2} \frac{e^{-2\rho|\mu_{n-1}|^\beta}}{1 + W(-\mu_{n-1})^{2c} + \xi_{n-1}^2} (C_n - \mu_{n-1} - \xi_{n-1})^2. \\ \times \mathbf{1}_{\{C_n - \mu_{n-1} \geq \xi_{n-1}\}} \frac{\tilde{p}^2(D_n - \mu_{n-1})}{\tilde{p}(D_n)\tilde{p}(D_n - 2\mu_{n-1})} \frac{\nabla \tilde{p}(D_n - 2\mu_{n-1})}{\tilde{p}(D_n - 2\mu_{n-1})}. \end{aligned} \quad (31)$$

In the above, (28) estimates the VaR, while (29) attempts to find the best variance reducer parameter for VaR estimation procedure. Similarly, (30) estimates the CVaR, while (29) attempts to find the best variance reducer parameter for CVaR estimation procedure.

Note on convergence. Since we approximated the true density $p(D)$ above using the policy, the convergence analysis of the above scheme is not straightforward. The difficult part is to establish that one can use the approximation $\tilde{p}(\cdot)$ in place of the true density $p(\cdot)$. Once this holds, then it can be shown that the tuple (η_n, μ_n) updated according to (29) and (31), converge to the optimal variance reducers (η^*, μ^*) , using arguments similar to that in Proposition 3.1 of Bardou et al. [2009]. (η^*, μ^*) minimize the convex functions

$$\begin{aligned} Q_1(\eta, \xi_\alpha^*) &:= \mathbb{E} \left[\mathbf{1}_{\{C^\theta(s^0) \geq \xi_\alpha^*\}} \frac{p(D)}{p(D-\eta)} \right] \text{ and} \\ Q_2(\mu, \xi_\alpha^*) &:= \mathbb{E} \left[(C^\theta(s^0) - \xi_\alpha^*)^2 \mathbf{1}_{\{C^\theta(s^0) \geq \xi_\alpha^*\}} \frac{p(D)}{p(D-\mu)} \right], \text{ where } \xi_\alpha^* \text{ is a VaR}_\alpha(C^\theta(s^0)). \end{aligned}$$

8 Comparison to previous work

In comparison to Borkar and Jain [2010] and Tamar et al. [2014], which are the most closely related contributions, we would like to point out the following:

- (i) The authors by Borkar and Jain [2010] develop an algorithm for a (finite horizon) CVaR constrained MDP, under a separability condition for the single-stage cost. On the other hand, without a separability condition, we devise policy gradient algorithms in a SSP setting and our algorithms are shown to converge as well; and
- (ii) The authors by Tamar et al. [2014] derive a likelihood estimate for the gradient of the CVaR of a random variable. However, they do not consider a risk-constrained SSP and instead optimize only CVaR. In contrast, we employ a convergent procedure for estimating CVaR that is motivated by a well-known convex optimization problem Rockafellar and Uryasev [2000] and then employ policy gradients for both the objective and constraints to find a locally risk-optimal policy.

9 Conclusions

In this paper, we proposed two novel algorithms to compute a risk-optimal policy in a stochastic shortest path problem. We used Conditional Value-at-Risk (CVaR) as a risk measure and this is motivated by applications in finance

and energy markets. Both the algorithms incorporated a CVaR estimation procedure along the lines of Bardou et al. [2009], which in turn is based on the well-known convex optimization representation by Rockafellar and Uryasev [2000]. For the purpose of finding a locally risk-optimal policy, our algorithms employed four tools: stochastic approximation, mini batches, policy gradients and importance sampling. In particular, stochastic approximation or mini-batches are used to approximate CVaR estimates/necessary gradients in the algorithms, while the gradients themselves are obtained using the likelihood ratio technique. Further, since CVaR is an expectation that conditions on the tail probability, to speed up CVaR estimation we incorporated an importance sampling procedure along the lines of Bardou et al. [2009]. We established asymptotic convergence of both the algorithms.

There are several future directions to be explored such as (i) obtaining finite-time bounds for our proposed algorithms, (ii) handling very large state spaces using function approximation, and (iii) applying our algorithms in practical contexts such as portfolio management in finance/energy sectors and revenue maximization in the re-insurance business.

References

- Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical finance*, 9(3):203–228, 1999.
- Yves F Atchade, Gersende Fort, and Eric Moulines. On stochastic proximal gradient algorithms. *arXiv preprint arXiv:1402.2365*, 2014.
- Olivier Bardou, Noufel Frikha, and Gilles Pages. Computing VaR and CVaR using stochastic approximation and adaptive unconstrained importance sampling. *Monte Carlo Methods and Applications*, 15(3):173–210, 2009.
- Peter L Bartlett and Jonathan Baxter. Infinite-horizon policy-gradient estimation. *arXiv preprint arXiv:1106.0665*, 2011.
- Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control, vol. II, 3rd edition*. Athena Scientific, 2007.
- V. Borkar. An actor-critic algorithm for constrained Markov decision processes. *Systems & Control Letters*, 54(3):207–213, 2005.
- V. Borkar. *Stochastic approximation: a dynamical systems viewpoint*. Cambridge University Press, 2008.
- Vivek Borkar and Rahul Jain. Risk-constrained Markov decision processes. In *IEEE Conference on Decision and Control (CDC)*, pages 2664–2669, 2010.
- Peter W Glynn. Likelihood ratio gradient estimation: an overview. In *Proceedings of the 19th conference on Winter simulation*, pages 366–375. ACM, 1987.
- H. Kushner and D. Clark. *Stochastic approximation methods for constrained and unconstrained systems*. Springer-Verlag, 1978.
- Vincent Lemaire and Gilles Pages. Unconstrained recursive importance sampling. *The Annals of Applied Probability*, 20(3):1029–1067, 2010.
- Shie Mannor and John Tsitsiklis. Mean-variance optimization in Markov decision processes. *arXiv preprint arXiv:1104.5601*, 2011.
- A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic theory*. Oxford University Press, 1995.
- L. A. Prashanth and M. Ghavamzadeh. Actor-critic algorithms for risk-sensitive MDPs. In *Proceedings of Advances in Neural Information Processing Systems 26*, pages 252–260, 2013.
- R Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.

- R. Sutton and A. Barto. *Reinforcement learning: An introduction*. MIT Press, 1998.
- A. Tamar and S. Mannor. Variance Adjusted Actor-Critic Algorithms. *arXiv preprint arXiv:1310.3697*, 2013.
- A. Tamar, D. Di Castro, and S. Mannor. Policy gradients with variance related risk criteria. In *Proceedings of the Twenty-Ninth International Conference on Machine Learning*, pages 387–396, 2012.
- Aviv Tamar, Yonatan Glassner, and Shie Mannor. Policy Gradients Beyond Expectations: Conditional Value-at-Risk. *arXiv preprint arXiv:1404.3862*, 2014.